Comparison of Routing Metrics for Static Multi-Hop Wireless Networks

Richard Draves Jitendra Padhye Brian Zill

Microsoft Research

{richdr, padhye, bzill}@microsoft.com

ABSTRACT

Routing protocols for wireless ad hoc networks have traditionally focused on finding paths with minimum hop count. However, such paths can include slow or lossy links, leading to poor throughput. A routing algorithm can select better paths by explicitly taking the quality of the wireless links into account. In this paper, we conduct a detailed, empirical evaluation of the performance of three link-quality metrics—ETX, per-hop RTT, and per-hop packet pair—and compare them against minimum hop count. We study these metrics using a DSR-based routing protocol running in a wireless testbed. We find that the ETX metric has the best performance when all nodes are stationary. We also find that the per-hop RTT and per-hop packet-pair metrics perform poorly due to self-interference. Interestingly, the hop-count metric outperforms all of the link-quality metrics in a scenario where the sender is mobile.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Networks

General Terms

Measurement, Performance, Experimentation

Keywords

Wireless multi-hop networks, Routing

1. INTRODUCTION

Routing in ad hoc wireless networks has been an active area of research for many years. Much of the original work in the area was motivated by mobile application environments, such as battlefield ad hoc networks. The primary focus in such environments is to provide scalable routing in the presence of mobile nodes. Recently, interesting commercial applications of multi-hop wireless networks have emerged. One example of such applications is “community wireless

networks” [27, 26, 16, 6]. In such networks, most of the nodes are either stationary or minimally mobile. The focus of routing algorithms in such networks is on improving the network capacity or the performance of individual transfers.

Most current ad hoc routing protocols select paths that minimize hop count [15, 24, 23, 22]. In static ad hoc wireless networks, minimal hop count paths can have poor performance because they tend to include wireless links between distant nodes. These long wireless links can be slow or lossy, leading to poor throughput [9]. A routing algorithm can select better paths by explicitly taking into account the quality of wireless links.

Researchers have proposed many metrics to measure wireless link quality, but to our knowledge, the relative performance of these metrics for the purpose of routing in static ad hoc wireless networks has not been investigated. In this paper, we study the performance of three link-quality metrics, and compare them against minimum hop-count routing. The first metric is called “Expected Transmission Count” (ETX). This metric is based on measuring the loss rate of broadcast packets between pairs of neighboring nodes [9]. The second metric is called “Per-hop Round Trip Time” (RTT). This metric is based on measuring the round trip delay seen by unicast probes between neighboring nodes [1]. The third metric is called “Per-hop Packet Pair Delay” (PktPair). This metric is based on measuring the delay between a pair of back-to-back probes to a neighboring node. We incorporated these metrics in an ad hoc routing protocol based on DSR. We deployed this protocol on a 23-node wireless testbed, and experimented with various traffic scenarios such as long-lived TCP flows, multiple simultaneous data transfers and simulated web traffic. We also considered a scenario involving some mobility.

The main contributions of the paper are the following: (i) We describe a routing protocol that incorporates the notion of link quality metrics. (ii) We present detailed experimental results to show that in scenarios with stationary nodes, the ETX metric outperforms hop-count although it uses longer paths. This is in contrast with the results in [9], in which the authors saw little or no gain from ETX in a DSR-based routing protocol. (iii) We show that the one-hop RTT and one-hop packet-pair metrics perform poorly, because their load-sensitivity leads to self-interference. (iv) We show that in a scenario involving a mobile sender, minimum hop-count routing performs considerably better than link-quality routing because the metrics do not react sufficiently quickly.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'04, Aug. 30–Sept. 3, 2004, Portland, Oregon, USA.

Copyright 2004 ACM 1-58113-862-8/04/0008 ...\$5.00.

2. LINK QUALITY METRICS

We consider three wireless link quality metrics in this paper. We also support minimum hop-count routing by defining a “HOP” metric. Each of these metrics represents a different notion of what constitutes good link quality. In Section 7, we will discuss other link quality metrics that we have not included in this study. The process of link discovery (i.e. neighbor discovery) is a separate issue, which we will discuss in in Section 3.

2.1 Hop Count (HOP)

This metric provides minimum hop-count routing. Link quality for this metric is a binary concept; either the link exists or it doesn’t.

The primary advantage of this metric is its simplicity. Once the topology is known, it is easy to compute and minimize the hop count between a source and a destination. Moreover, computing the hop count requires no additional measurements, unlike the other metrics we will describe in this section.

The primary disadvantage of this metric is that it does not take packet loss or bandwidth into account. It has been shown [9] that a route that minimizes the hop count does not necessarily maximize the throughput of a flow. For example, a two-hop path over reliable or fast links can exhibit better performance than a one-hop path over a lossy or slow link. The HOP metric, however, will prefer the one-hop path.

2.2 Per-hop Round Trip Time (RTT)

This metric is based on measuring the round trip delay seen by unicast probes between neighboring nodes. Adya et al. [1] proposed this metric. To calculate RTT, a node sends a probe packet carrying a timestamp to each of its neighbors every 500 milliseconds. Each neighbor immediately responds to the probe with a probe acknowledgment, echoing the timestamp. This enables the sending node to measure round trip time to each of its neighbors. The node keeps an exponentially weighted moving average of the RTT samples to each of its neighbors. Our implementation gives 10% weight to the current sample while calculating the average. If a probe or a probe response packet is lost, the average is increased by 20% to reflect this loss. Similar penalty is taken if loss of a data packet is detected on the link. We also increase the average if we detect a loss of data packet. The routing algorithm selects the path with the least total sum of RTTs.

The RTT metric measures several facets of link quality. First, if either the node or the neighbor is busy, the probe or the probe-ack packet will experience queuing delay, resulting in high RTT. Second, as shown in [1], if other nodes in the vicinity are busy, the probe or the probe-ack packet will experience delays due to channel contention, again resulting in high RTT. Third, if link between the nodes is lossy, the 802.11 ARQ mechanism may have to retransmit the probe or the probe-ack packet several times to get it delivered correctly. This increases the RTT along that hop. Finally, if despite the ARQ mechanism, a probe or a probe-ack packet is lost, the sender node detect the loss, and increases the moving average as described earlier. In short, the RTT metric is designed to avoid highly loaded or lossy links.

Since RTT is a load-dependent metric, it can lead to route instability. This is a well-known problem in wired networks [18, 2]. We call this phenomenon self-interference.

The route instability is made worse by the fact that due to limitations of our implementation, we are unable to insert the probe packets at the head of the queue maintained by the driver. This queuing delay significantly distorts the RTT value on that hop. The authors of [1] have also reported this problem. This metric has other disadvantages as well. First, there is the overhead of measuring the round trip time. We reduce this overhead by using small probe packets (137 bytes). Second, the metric doesn’t explicitly take link data rate into account. We may be able to take impact of link data rate into account by using larger probe packets. However, larger probes would impose an even greater measurement overhead. Finally, this measurement technique requires that every pair of neighboring nodes probe each other. Thus, the technique might not scale to dense networks.

2.3 Per-hop Packet Pair Delay (PktPair)

This metric is based on measuring the delay between a pair of back-to-back probes to a neighboring node. It is designed to correct the problem of distortion of RTT measurement due to queuing delays. The packet-pair technique is well-known in the world of wired networks [17].

To calculate this metric, a node sends two probe packets back-to-back to each neighbor every 2 seconds. The first probe packet is small, and the next one is large. The neighbor calculates the delay between the receipt of the first and the second packets. It then reports this delay back to the sending node. The sender maintains a exponentially weighted moving average of these delays for each of its neighbors. The objective of the routing algorithm is to minimize the sum of these delays.

Like the RTT metric, this metric also measures several facets of link quality. If, due to high loss rate, the second probe packet requires retransmissions by 802.11 ARQ, the delay measured by the neighbor will increase. If the link from the node to its neighbor has low bandwidth, the second packet will take more time to traverse the link, which will result in increased delay. If there is traffic in the vicinity of this hop, it will also result in increased delay, since the probe packets have to contend for the channel.

The primary advantage of this metric over RTT is that it isn’t affected by queuing delays at the sending node, since both packets in a pair will be delayed equally. In addition, using a larger packet for the second probe makes the metric more sensitive to the link bandwidth than the RTT metric.

This metric has several disadvantages. First, it is subject to overheads even greater than those of the RTT metric, since two packets are sent to each neighbor, and the second packet is larger. Second, we discovered that the metric is not completely immune to the phenomenon of self-interference. To understand self-interference for packet-pair measurements, consider three wireless nodes A, B, and C forming a two-hop chain topology. Assume that A is sending data to C via B. If a queue builds up on the link from A to B, the PktPair measurements on that link won’t be affected because both the probe packets would be delayed equally. Now consider the link from B to C. Node B can not simultaneously receive a packet from A and send a probe to C. This means that the probe packet is contending with the data packet for the wireless channel. This increases the metric from B to C, increasing the total metric along the path from A to C. However, this self-interference is less severe than that experienced by RTT.

IPv4	IPv6	IPX, etc.
Mesh Connectivity Layer (with LQSR)		
Ethernet	802.11	802.16, etc.

Figure 1: Our architecture multiplexes multiple physical links into a single virtual link.

2.4 Expected Transmission Count (ETX)

This metric estimates the number of retransmissions needed to send unicast packets by measuring the loss rate of broadcast packets between pairs of neighboring nodes. De Couto et al. [9] proposed ETX. To compute ETX, each node broadcasts a probe packet every second. The probe contains the count of probes received from each neighboring node in the previous 10 seconds. Based on these probes, a node can calculate the loss rate of probes on the links to and from its neighbors. Since the 802.11 MAC does not retransmit broadcast packets, these counts allow the sender to estimate the number of times the 802.11 ARQ mechanism will retransmit a unicast packet.

To illustrate this, consider two nodes A and B. Assume that node A has received 8 probe packets from B in the previous 10 seconds, and in the last probe packet, B reported that it had received 9 probe packets from A in the previous 10 seconds. Thus, the loss rate of packets from A to B is 0.1, while the loss rate of packets from B to A is 0.2. A successful unicast data transfer in 802.11 involves sending the data packet and receiving a link-layer acknowledgment from the receiver. Thus, the probability that the data packet will be successfully transmitted from A to B in a single attempt is $(1 - 0.1) \times (1 - 0.2) = 0.72$. If either the data or the ack is lost, the 802.11 ARQ mechanism will retransmit the packet. If we assume that losses are independent, the expected number of retransmissions before the packet is successfully delivered is $1/0.72 = 1.39$. This is the value of the ETX metric for the link from A to B. The routing protocol finds a path that minimizes the sum of the expected number of retransmissions.

Node A calculates a new ETX value for the link from A to B every time it receives a probe from B. In our implementation of the ETX metric, the node maintains an exponentially weighted moving average of ETX samples. There is no question of taking 20% penalty for lost probe packets. Penalty is taken only upon loss of a data packet.

ETX has several advantages. Since each node broadcasts the probe packets instead of unicasting them, the probing overhead is substantially reduced. The metric suffers little from self-interference since we are not measuring delays.

The main disadvantage of this metric is that since broadcast probe packets are small, and are sent at the lowest possible data rate (6Mbps in case of 802.11a), they may not experience the same loss rate as data packets sent at higher rates. Moreover, the metric does not directly account for link load or data rate. A heavily loaded link may have very low loss rate, and two links with different data rates may have the same loss rate.

3. AD HOC ROUTING ARCHITECTURE

We implement ad hoc routing and link-quality measurement in a module that we call the Mesh Connectivity Layer (MCL). Architecturally, MCL is a loadable Windows driver.

It implements a virtual network adapter, so that to the rest of the system the ad hoc network appears as an additional (virtual) network link. MCL routes using a modified version of DSR [15] that we call Link-Quality Source Routing (LQSR). We have modified DSR extensively to improve its behavior, most significantly to support link-quality metrics. In this section, we briefly review our architecture and implementation to provide background for understanding the performance results. More architectural and implementation details are available in [10].

The MCL driver implements an interposition layer between layer 2 (the link layer) and layer 3 (the network layer). To higher-layer software, MCL appears to be just another ethernet link, albeit a virtual link. To lower-layer software, MCL appears to be just another protocol running over the physical link. See Figure 1 for a diagram.

This design has two significant advantages. First, higher-layer software runs unmodified over the ad hoc network. In our testbed, we run both IPv4 and IPv6 over the ad hoc network. No modifications to either network stack were required. Second, the ad hoc routing runs over heterogeneous link layers. Our current implementation supports ethernet-like physical link layers (eg 802.11 and 802.3). The virtual MCL network adapter can multiplex several physical network adapters, so the ad hoc network can extend across heterogeneous physical links.

In the simple configuration shown in Figure 1, the MCL driver binds to all the physical adapters and IP binds only to the MCL virtual adapter. This avoids multi-homing at the IP layer. However other configurations are also possible. In our testbed deployment, the nodes have both an 802.11 adapter for the ad hoc network and an ethernet adapter for management and diagnosis. We configure MCL to bind only to the 802.11 adapter. The IP stack binds to both MCL and the ethernet adapter. Hence the mesh nodes are multi-homed at the IP layer, so they have both a mesh IP address and a management IP address. We prevent MCL from binding to the management ethernet adapter, so the ad hoc routing does not discover the ethernet as a high-quality single-hop link between all mesh nodes.

The MCL adapter has its own 48-bit virtual ethernet address, distinct from the layer-2 addresses of the underlying physical adapters. The mesh network functions just like an ethernet, except that it has a smaller MTU. To allow room for the LQSR headers, it exposes a 1280-byte MTU instead of the normal 1500-byte ethernet MTU. Our 802.11 drivers do not support the maximum 2346-byte 802.11 frame size.

The MCL driver implements a version of DSR that we call Link-Quality Source Routing (LQSR). LQSR implements all the basic DSR functionality, including Route Discovery (Route Request and Route Reply messages) and Route Maintenance (Route Error messages). LQSR uses a link cache instead of a route cache, so fundamentally it is a link-state routing protocol like OSPF [20]. The primary changes in LQSR versus DSR relate to its implementation at layer 2.5 instead of layer 3 and its support for link-quality metrics.

Due to the layer 2.5 architecture, LQSR uses 48-bit virtual ethernet addresses. All LQSR headers, including Source Route, Route Request, Route Reply, and Route Error, use 48-bit virtual addresses instead of 32-bit IP addresses.

We have modified DSR in several ways to support routing according to link-quality metrics. These include modifications to Route Discovery and Route Maintenance plus new

mechanisms for Metric Maintenance. Our design does not assume that the link-quality metric is symmetric.

First, LQSR Route Discovery supports link metrics. When a node receives a Route Request and appends its own address to the route in the Route Request, it also appends the metric for the link over which the packet arrived. When a node sends a Route Reply, the reply carries back the complete list of link metrics for the route.

Once Route Discovery populates a node's link cache, the cached link metrics must be kept reasonably up-to-date for the node's routing to remain accurate. In Section 5.2 we show that link metrics do vary considerably, even when nodes are not mobile. LQSR tackles this with two separate Metric Maintenance mechanisms.

LQSR uses a reactive mechanism to maintain the metrics for the links which it is actively using. When a node sends a source-routed packet, each intermediate node updates the source route with the current metric for the next (outgoing) link. This carries up-to-date link metrics forward with the data. To get the link metrics back to the source of the packet flow (where they are needed for the routing computation), we have the recipient of a source-routed data packet send a gratuitous Route Reply back to the source, conveying the up-to-date link metrics from the arriving Source Route. This gratuitous Route Reply is delayed up to one second waiting for a piggy-backing opportunity.

LQSR uses a proactive background mechanism to maintain the metrics for all links. Occasionally each LQSR node send a Link Info message. The Link Info carries current metrics for each link from the originating node. The Link Info is piggy-backed on a Route Request, so it floods throughout the neighborhood of the node. LQSR piggy-backs Link Info messages on all Route Requests when possible.

The link metric support also affects Route Maintenance. When Route Maintenance notices that a link is not functional (because a requested Ack has not been received), it penalizes the link's metric and sends a Route Error. The Route Error carries the link's updated metric back to the source of the packet.

Our LQSR implementation includes the usual DSR conceptual data structures. These include a Send Buffer, for buffering packets while performing Route Discovery; a Maintenance Buffer, for buffering packets while performing Route Maintenance; and a Route Request table, for suppressing duplicate Route Requests. The LQSR implementation of Route Discovery omits some optimizations that are not worthwhile in our environment. In practice, Route Discovery is almost never required in our testbed so we have not optimized it. In particular, LQSR nodes do not reply to Route Requests from their link cache. Only the target of a Route Request sends a Route Reply. Furthermore, nodes do not send Route Requests with a hop limit to restrict their propagation. Route Requests always flood throughout the ad hoc network. Nodes do cache information from overheard Route Requests.

The Windows 802.11 drivers do not support promiscuous mode and they do not indicate whether a packet was successfully transmitted. Hence our implementation of Route Maintenance uses explicit acknowledgments instead of passive acknowledgments or link-layer acknowledgments. Every source-routed packet carries an Ack Request option. A node expects an Ack from the next hop within 500ms. The Ack options are delayed briefly (up to 80ms) so that they may be

piggy-backed on other packets flowing in the reverse direction. Also later Acks squash (replace) earlier Acks that are waiting for transmission. As a result of these techniques, the acknowledgment mechanism does not add significant byte or packet overhead.

The LQSR implementation of Route Maintenance also omits some optimizations. We do not implement "Automatic Route Shortening," because of the lack of promiscuous mode. We also do not implement "Increased Spreading of Route Error Messages". This is not important because LQSR will not reply to a Route Request from (possibly stale) cached data. When LQSR Route Maintenance detects a broken link, it does not remove from the transmit queue other packets waiting to be sent over the broken link, since Windows drivers do not provide access to the transmit queue. However, LQSR nodes do learn from Route Error messages that they forward.

LQSR supports a form of DSR's "Packet Salvaging" or retransmission. Salvaging allows a node to try a different route when it is forwarding a source-routed packet and discovers that the next hop is not reachable. The acknowledgment mechanism does not allow every packet to be salvaged because it is primarily designed to detect when links fail. When sending a packet over a link, if the link has recently (within 250ms) been confirmed to be functional, we request an Ack as usual but we do not buffer the packet for possible salvaging. This design allows for salvaging the first packets in a new connection and salvaging infrequent connection-less communication, but relies on transport-layer retransmission for active connections. In our experience, packets traversing "cold" routes are more vulnerable to loss from stale routes and benefit from the retransmission provided by salvaging.

We have not yet implemented the DSR "Flow State" optimization, which uses soft-state to replace a full source route with a small flow identifier. We intend to implement it in the future. Our Link Cache implementation does not use the Link-MaxLife algorithm [15] to timeout links. We found that Link-MaxLife produced inordinate churn in the link cache. Instead, we use an infinite metric value to denote broken links in the cache. We garbage-collect dead links in the cache after a day.

4. TESTBED

The experimental data reported in this paper are the results of measurements we have taken on a 23-node wireless testbed. Our testbed is located on one floor of a fairly typical office building, with the nodes placed in offices, conference rooms and labs. Unlike wireless-friendly cubicle environments, our building has rooms with floor-to-ceiling walls and solid wood doors. With the exception of one additional laptop used in the mobility experiments, the nodes are located in fixed locations and did not move during testing. The node density was deliberately kept high enough to enable a wide variety of multi-hop path choices. See Figure 2.

The nodes are primarily laptop PCs with Intel Pentium II processors with clock rates from 233 to 300 MHz, but also included a couple slightly faster laptops as well as two desktop machines. All of the nodes run Microsoft Windows XP. The TCP stack included with XP supports the SACK option by default, and we left it enabled for all of our TCP experiments. All of our experiments were conducted over IPv4 using statically assigned addresses. Each node has an 802.11a PCCARD radio. We used the default configuration

Figure 2: Our testbed consists of 23 nodes placed in fixed locations inside an office building. Four different models of 802.11a wireless cards were used. The six shaded nodes were used as endpoints for a subset of the experiments (see section 5.3). The mobile node walking path shows the route taken during the mobility experiments (see section 6).

for the radios, except for configuring ad hoc mode and channel 36 (5.18 GHz). In particular, the cards all performed auto-speed selection. There are no other 802.11a users in our building.

We use four different types of cards in our testbed: 11 Proxim ORiNOCO ComboCard Gold, 7 NetGear WAG 511, 4 NetGear WAB 501, and 1 Proxim Harmony. While we performed no formal testing of radio ranges, we observed that some cards exhibited noticeably better range than others. The Proxim ORiNOCOs had the worst range of the cards we used in the testbed. The NetGear WAG 511s and WAB 501s exhibited range comparable to each other, and somewhere between the two Proxim cards. The Proxim Harmony had the best range of the cards we tried.

5. RESULTS

In this section, we describe the results of our experiments. The section is organized as follows. First, we present measurements that characterize our testbed. These include a study of the overhead imposed by our routing software, and a characterization of the variability of wireless links in our testbed. Then, we present experiments that compare the four routing metrics under various type of traffic.

5.1 LQSR Overhead

Like any routing protocol, LQSR incurs a certain amount of overhead. First, it adds additional traffic in the form of routing updates, probes, etc. Second, it has the overhead of carrying the source route and other fields in each packet. Third, all nodes along the path of a data flow sign each packet using HMAC-SHA1 and regenerate the hash to reflect the changes in the LQSR headers when forwarding the packet. Also, the end nodes encrypt or decrypt the payload data using AES-128, since our sysadmin insisted upon better protection than WEP. The cryptographic overhead can be significant given the slow CPU speeds of the nodes.

We conducted the following experiment measures the overhead of LQSR. We equipped four laptops (named A, B, C and D), with a Proxim ORiNOCO card each. The laptops were placed in close proximity of each other. All the links between the machines were operating at the maximum data rate (nominally 54Mbps).

To establish a baseline for our measurements, we set up static IP routes between these machines to form a chain topology. In other words, a packet from A to D was forwarded via B and C, and a packet from D to A was forwarded via C and B. We measured the throughput of long-lived TCP flows from A to B (a 1-hop path), from A to C (a 2-hop path) and from A to D (a 3-hop path). At any time, only one TCP flow was active. These throughputs form our baseline.

Next, we deleted the static IP routes and started LQSR on each machine. LQSR allows the user to set up static routes that override the routes discovered by the routing protocol. We set up static routes to form the same chain topology described earlier. Note that LQSR continues to send its normal control packets and headers, but routes discovered through this process are ignored in favor of the static routes supplied by the user. We once again measured the throughput of long-lived TCP flows on 1, 2 and 3 hop paths. Finally, we turned off all cryptographic functionality and measured the throughput over 1, 2 and 3 hop paths again.

The results of these experiments are shown in Figure 3. Each bar represents the average of 5 TCP connections. The variation between runs is negligible. The first thing to note is that, as one would expect, the throughput falls linearly with the number of hops due to inter-hop interference. LQSR's overhead is most evident on 1-hop paths. The throughput reduction due to LQSR, compared to static IP routing is over 38%. However, when cryptographic functionality is turned off, the throughput reduction is only 13%. Thus, we can conclude that the LQSR overhead is largely due to cryptography, which implies that the CPU is the bottleneck. The amount of overhead imposed by LQSR decreases as the path length increases. This is because channel contention between successive hops is the dominant cause of throughput reduction. At these reduced data rates, the CPU can easily handle the cryptographic overhead. The remaining overhead is due to the additional network traffic and headers carried in each packet. Note that the amount of control traffic varies depending on the number of neighbors that a node has. This is especially true for the RTT and PktPair metrics. We believe that this variation is not significant.

Figure 3: MCL overhead is not significant on multi-hop paths.

5.2 Link Variability in the Testbed

In our testbed, we allow the radios to dynamically select their own data rates. Thus, different links in our testbed have different bandwidths. To characterize this variability in link quality, we conducted the following experiment. Recall the PktPair metric collects a sample of the amount of time required to transmit a probe packet every 2 seconds on each link. We modified the implementation of the PktPair metric to keep track of the minimum sample out of every successive 50 samples (i.e. minimum of samples 1-50, then minimum of samples 51-100 etc.). We divide the size of the second packet by this minimum. The resulting number is an indication of the bandwidth of that link during that 100 second period. In controlled experiments, we verified that this approach approximates the raw link bandwidth. We gathered these samples from all links for a period of 14 hours. Thus, for each link there were a total of $14 \times 60 \times 60 \div 100 = 504$ bandwidth samples. There was no traffic on the testbed during this time. We discard any intervals in which the calculated bandwidth is more than 36Mbps, which is the highest data rate that we actually see in the testbed. This resulted in removal of 3.83% of all bandwidth samples. Still, the resulting number is not an exact measure of the available bandwidth, since it is difficult to correctly account for all link-layer overhead. However, we believe that the number is a good (but rough) indication of the link bandwidth during the 100 second period.

Of the $23 \times 22 = 506$ total possible links, only 183 links had non-zero average bandwidth, where the average was computed across all samples gathered over 14 hours. We found that bandwidths of certain links varied significantly over time, while other links it was relatively stable. Examples of two such links appear in Figures 4 and 5. In the first case, we see that the bandwidth is relatively stable over the duration of the experiment. In the second case, however, the bandwidth is much more variable. Since the quality of links in our testbed varies over time, we are careful to repeat our experiments at different times.

In Figure 6 we compare the bandwidth on the forward and reverse direction of a link. To do this, we consider all possible unordered node pairs. The number of such pairs is $23 \times 22 \div 2 = 253$. Each pair corresponds to two directional links. Each of these two links has its own average bandwidth. Thus, each pair has two bandwidths associated with it. Out of the 253 possible node pairs, 90 node pairs had non-zero average bandwidth in both forward and reverse directions. In Figure 6, we plot a point for each such pair. The X-coordinate of the pair represents the link with the larger bandwidth. The existence of several points below the

diagonal line implies that there are several pairs for which the forward and the reverse bandwidths differ significantly. In fact, in 47 node pairs, the forward and the reverse bandwidths differ by more than 25%.

5.3 Impact on Long Lived TCP Flows

Having characterized the overhead of our routing software, and the quality of links in our testbed, we can now discuss how various routing metrics perform in our testbed. We begin by discussing the impact of routing metrics on the performance of long-lived TCP connections. In today's Internet, TCP carries most of the traffic, and most of the bytes are carried as part of long-lived TCP flows [13]. It is reasonable to expect that similar types of traffic will be present on community networks, such as [6, 27]. Therefore, it is important to examine the impact of routing metrics on the performance of long-lived TCP flows.

We start the performance comparison with a simple experiment. We carried out a TCP transfer between each unique sender-destination pair. There are 23 nodes in our testbed, so a total of $23 \times 22 = 506$ TCP transfers were carried out. Each TCP transfer lasted for 3 minutes, and transferred as much data as it could. On the best one-hop path in our testbed a 3 minute connection will transfer over 125MB of data. Such large TCP transfers ensure repeatability of results. We had previously determined empirically that TCP connections of 1 minute duration were of sufficient length to overcome startup effects and give reproducible results. We used 3 minute transfers in these experiments, to be conservative. We chose to fix the transfer duration, instead of the amount of data transferred, to keep the running time of each experiment predictable. Only one TCP transfer was active at any time. The total time required for the experiment was just over 25 hours. We repeated the experiment for each metric.

In Figure 7 we show the median throughput of the 506 TCP transfers for each metric. We choose median to represent the data instead of the mean because the distribution (which includes transfers of varying path lengths) is quite skewed. The height error bars represent the *semi-inter quartile range* (SIQR), which is defined as half the difference between 25th and 75th percentile of the data. SIQR is the recommended measure of dispersion when the central tendency of the data is represented by the median [14]. Since each connection is run between a different pair of nodes the relatively large error bars indicate that we observe a wide range of throughputs across all the pairs. The median throughput using the HOP metric is 1100Kbps, while the median throughput using the ETX metric is 1357Kbps. This represents an improvement of 23.1%.

In contrast, De Couto et al. [9] observed almost no improvement for ETX in their DSR experiments. There are several possible explanations for this. First, they used UDP instead of TCP. A bad path will have more impact on the throughput of a TCP connection (due to window backoffs, timeouts etc.) than on the throughput of a UDP connection. Hence, TCP amplifies the negative impact of poor route selection. Second, in their testbed the radios were set to their lowest sending rate of 1Mbps, whereas we allow the radios to set transmit rates automatically (auto-rate). We believe links with lower loss rates also tend to have higher data rates, further amplifying ETX's improvement. Third, our testbed has 6-7 hop diameter whereas their testbed has

Figure 4: The bandwidth of the link from node 65 to node 100 is relatively stable.

Figure 5: The bandwidth of the link from node 68 to node 66 varies over time.

Figure 6: For many node pairs, the bandwidths of the two links between the nodes is unequal.

Figure 7: All pairs: Median throughput of TCP transfer.

Figure 8: All Pairs: Median number of paths per TCP transfer.

Figure 9: All pairs: Median path length of a TCP transfer.

Figure 10: All Pairs: Comparison of HOP and ETX path lengths. ETX consistently uses longer paths than HOP.

a 5-hop diameter [8]. As we discuss below, ETX’s improvement over HOP is more pronounced at longer path lengths.

The RTT metric gives the worst performance among the four metrics. This is due to the phenomenon of self-interference that we previously noted in Section 2.2. The phenomenon manifests itself in the number of paths taken by the connection. At the beginning the connection uses a certain path. However, due to self-interference, the metric on this path soon rises. The connection then chooses another path. This is illustrated in Figure 8. The graph shows the median number of paths taken by a connection. The RTT metric uses far more paths per connection than other metrics. The HOP metric uses the least number of paths per connection - the median is just 1.

The PktPair metric performs better than RTT, but worse than both HOP and ETX. This is again due to the phenomenon of self-interference. While the RTT metric suffers from self-interference on all hops along the path, the PktPair metric eliminates the self-interference problem on the first hop. The impact of this can be seen in the median number of paths (12) tried by a connection using the PktPair met-

ric. This number is lower than median using RTT (20.5), but substantially higher than HOP and ETX.

Note that the ETX metric also uses several paths per connection: the median is 4. This is because for a given node pair, multiple paths that are essentially equivalent can exist between them. There are several such node pairs in our testbed. Small fluctuations in the metric values of equivalent paths can make ETX choose one path over another. We plan to investigate route damping strategies to alleviate this problem.

The self-interference, and consequent route flapping experienced the RTT metric has also been observed in wired networks [18, 2]. In [18], the authors propose to solve the problem by converting the RTT to utilization, and normalizing the resulting value for use as a route metric. In [2], the authors propose to use hysteresis to alleviate route flapping. We are currently investigating these techniques further. Our initial results show that hysteresis may reduce the severity of the problem, but not significantly so.

5.4 Impact of Path Length

The HOP metric produces significantly shorter paths than the three other metrics. This is illustrated in Figure 9. The bar chart shows the median across all 506 TCP transfers of the average path length of each transfer. To calculate the average path length of a TCP transfer, we keep track of the paths taken by all the data-carrying packets in the transfer. We calculate the average path length by weighting the length of each unique path by the number of packets that took that path. The error bars represent SIQR. The HOP metric has the shortest median path length (2), followed by ETX (3.01), RTT (3.43) and PktPair (3.58).

We now look at ETX and HOP path lengths in more detail. In Figure 10, we plot the average path length of each TCP transfer using HOP versus the average path length

Figure 11: All Pairs: Throughput as a function of path length under HOP. The metric does a poor job of selecting multi-hop paths.

Figure 12: All Pairs: Throughput as a function of path length under ETX. The metric does a better job of selecting multi-hop paths.

Figure 13: All Pairs: Throughput as a function of path length under RTT. The metric does a poor job of selecting even one hop paths.

Figure 14: All Pairs: Throughput as a function of path length under PktPair. The metric finds good one-hop paths, but poor multi-hop paths.

using ETX. Again, the ETX metric produces significantly longer paths than the HOP metric. The testbed diameter is 7 hops using ETX and 6 hops using HOP.

We also examined the impact of average path length on TCP throughput. In Figure 11 we plot the throughput of a TCP connection against its path length using HOP while in Figure 12 we plot the equivalent data for ETX. First, note that as one would expect, longer paths produce lower throughputs because channel contention keeps more than one link from being active. Second, note that ETX’s higher median throughput derives more from avoiding lower throughputs than from achieving higher throughputs. Third, ETX does especially well at longer path lengths. The ETX plot is flat from around 5 through 7 hops, possibly indicating that links at opposite ends of the testbed do not interfere. Fourth, ETX avoids poor one-hop paths whereas HOP blithely uses them.

We now look at the performance of RTT and PktPair in more detail. In Figure 13 we plot TCP throughput versus average path length for RTT while in Figure 14 we plot the data for PktPair. RTT’s self-interference is clearly evident in the low throughputs and high number of connections with average path lengths between 1 and 2 hops. With RTT, even 1-hop paths are not stable. In contrast, with PktPair the 1-hop paths look good (equivalent to ETX in Figure 12) but self-interference is evident starting at 2 hops.

5.5 Variability of TCP Throughput

To measure the impact of routing metrics on the variability of TCP throughput, we carry out the following experi-

ment. We select 6 nodes on the periphery of the testbed, as shown in Figure 2. Each of the 6 nodes then carried out a 3-minute TCP transfer to the remaining 5 nodes. The TCP transfers were set sufficiently apart to ensure that no more than one TCP transfer will be active at any time. There is no other traffic in the testbed. We repeated this experiment 10 times. Thus, there were a total of $6 \times 5 \times 10 = 300$ TCP transfers. Since each transfer takes 3 minutes, the experiment lasts for over 15 hours.

In Figure 15 we show the median throughput of the 300 TCP transfers using each metric. As before, the error bars represent SIQR. Once again, we see that the RTT metric is the worst performer, and the ETX metric outperforms the other three metrics by a wide margin.

The median throughput using the ETX metric is 1133Kbps, while the median throughput using the HOP metric is 807.5. This represents a gain of 40.3%. This gain is higher than the 23.15% obtained in the previous experiment because these machines are on the periphery of the network, and thus, the paths between them tend to be longer. As we have noted in Section 5.4, the ETX metric tends to perform better than HOP on longer paths. The higher median path lengths substantially degrades the performance of RTT and PktPair, compared to their performance shown in Figure 7.

The HOP metric selects the shortest path between a pair of nodes. If multiple shortest paths are available, the metric simply chooses the first one it finds. This introduces a certain amount of randomness in the performance of the HOP metric. If multiple TCP transfers are carried out between a given pair of nodes, the HOP metric may select different paths for each transfer. The ETX metric, on the other hand, selects “good” links. This means that it tends to choose the same path between a pair of nodes, as long as the link qualities do not change drastically. Thus, if several TCP transfers are carried out between the same pair of nodes at different times, they should yield similar throughput using ETX, while the throughput under HOP will be more variable. This fact is illustrated in Figure 16.

The figure uses coefficient of variation (CoV) as a measure of variability. CoV is defined as standard deviation divided by mean. There is one point in the figure for each of the 30 source-destination pairs. The X-coordinate represents CoV of the throughput of 10 TCP transfers conducted between a given source-destination pair using ETX, and the Y coordinate represents the CoV using HOP. The CoV values are

Figure 15: 30 Pairs: Median throughput with different metrics.

significantly lower with ETX. Note that a single point lies well below the diagonal line, indicating that HOP provided more stable throughput than ETX. This point represents TCP transfers from node 23 to node 10. We are currently investigating these transfers further. It is interesting to note that for the reverse transfers, i.e. from node 10 to node 23, ETX provides lower CoV than than HOP.

5.6 Multiple Simultaneous TCP Transfers

In the experiments described in the previous section, only one TCP connection was active at any time. This is unlikely to be the case in a real network. In this section, we compare the performance of ETX, HOP and PktPair for multiple simultaneous TCP connections. We do not consider RTT since its performance is poor even with a single TCP connection.

We use the same set of 6 peripheral nodes shown in Figure 2. We establish 10 TCP connections between each distinct pair of nodes. Thus, there are a total of $6 \times 5 \times 10 = 300$ possible TCP connections. Each TCP connection lasts for 3 minutes. The order in which the connections are established is randomized. The wait time between the start of two successive connections determines the number of simultaneously active connections. For example, if the wait time between starting consecutive connections is 90 seconds, then two connections will be active simultaneously. We repeat the experiment for various numbers of simultaneously active connections.

For each experiment we calculate the median throughput of the 300 connections, and multiply it by the number of simultaneously active connections. We call this product the Multiplied Median Throughput (MMT). MMT should increase with the number of simultaneous connections, until the load becomes too high for the network to carry.

In Figure 17 we plot MMT against the number of simultaneous active connections. The figure shows that the performance of the PktPair metric gets significantly worse as the number of simultaneous connections increase. This is because the self-interference problem gets worse with increasing load. In the case of ETX, the MMT increases to a peak at 5 simultaneous connections. The MMT growth is significantly less than linear because there is not much parallelism in our testbed (many links interfere with each other) and the increase that we are seeing is partly because a single TCP connection does not fully utilize the end-to-end path. We believe the MMT falls beyond 5 simultaneous connections due to several factors, including 802.11 MAC inefficiencies and instability in the ETX metric under very high load. The MMT using HOP deteriorates much faster

Figure 16: 30 Pairs: The lower CoVs under ETX indicate that ETX chooses stable links.

Figure 17: Throughputs with multiple simultaneous TCP connections.

than it does with ETX. As discussed in Section 5.8, at higher loads HOP performance drops because link-failure detection becomes less effective.

5.7 Web-like TCP Transfers

Web traffic constitutes a significant portion of the total Internet traffic today. It is reasonable to assume that web traffic will also be a significant portion of traffic in wireless meshes such as the MIT Roofnet [26]. The web traffic is characterized by the heavy-tailed distribution of flow sizes: most transfers are small, but there are some very large transfers [21]. Thus, it is important to examine the performance of web traffic under various routing metrics.

To conduct this experiment, we set up a web server on host 128. The six peripheral nodes served as web clients. The web traffic was generated using Surge [5]. The Surge software has two main parts, a file generator and a request generator. The file generator generate files of varying sizes that are placed on the web server. The Surge request generator models a web user that fetches these files. The file generator and the request generator offer a large number of parameters to customize file size distribution and user behaviors. We ran Surge with its default parameter settings, which have been derived through modeling of empirical data [5].

Each Surge client modeled a single user running HTTP 1.1. Each user session lasted for 40 minutes, divided in four slots of 10 minutes each. Each user fetched over 1300 files from the web server. The smallest file fetched was 77 bytes long, while the largest was 700KB. We chose to have only one client active at any time, to allow us to study the behavior of each client in detail. We measure the latency or each object: the amount of time elapsed between the request for an object, and the completion of its receipt. Note that we are ignoring any rendering delays.

Figure 18: Median latency for all files fetched

Figure 19: Median latency for files smaller than 1KB

Figure 20: Median latency for files larger than 8KB

In Figure 18, we plot the median latencies experienced by each client. It can be seen that ETX reduces the latencies observed by clients that are further away from the web server. This is consistent with our earlier finding ETX tends to perform better than HOP on longer paths. For host 23, 100 and 20, the median latency under ETX is almost 20% lower than the median latency under HOP. These hosts are relatively further away from the webserver running on host 128. On the other hand, for host 1, the median latency under HOP is lower by 20%. Host 1 is just one hop away from the web server. These results are consistent with the results in Section 5.4: on longer paths, ETX performs better than HOP, but on one-hop paths, the HOP metric sometimes performs better.

To study whether the impact of ETX is limited to large transfers we studied the median response times for small objects, i.e. files that are less than 1KB in size and large objects, i.e., those over 8KB in size. These medians are shown in Figures 19 and 20, respectively. The benefit of ETX is indeed more evident in case of larger transfers. However, ETX also reduces the latency of small transfers by significant proportion. This is particularly interesting as the data sent from the server to client in such small transfers fits inside a single TCP packet. It is clear that even for such short transfers, the paths selected by ETX are better.

5.8 Discussion

We conclude from our results that load-sensitivity is the primary factor determining the relative performance of the three metrics. The RTT metric is the most sensitive to load; it suffers from self-interference even on one-hop paths and has the worst performance. The PktPair metric is not affected by load generated by the probing node, but it is sensitive to other load on the channel. This causes self-interference on multi-hop paths and degrades performance. The ETX metric has the least sensitivity to load and it performs the best.

Our experience with HOP leads us to believe that its performance is very sensitive to competing factors controlling the presence of poor links in the link cache. Consider the evolution of a route during a long data transfer. When the transfer starts, the link cache contains a fairly complete topology, including many poor links. The shortest-path Dijkstra computation picks a route that probably includes poor (lossy or slow) links. Then as the data transfer proceeds, Route Maintenance detects link failures and sends Route Error messages, causing the failed link to be removed from link caches. Since poor links suffer link failures more frequently than good links, over time the route tends to im-

Figure 21: Median Throughput of 45 1-minute TCP transfers with mobile sender using HOP and ETX.

prove. However this process can go too far: if too many links are removed, the route can get longer and longer until finally there is no route left and the node performs Route Discovery again. On the other hand, a background load of unrelated traffic in the network tends to repopulate the link cache with good and bad links, because of the caching of link existence from overheard or forwarded packets. The competition between these two factors, one removing links from the link cache and the other adding links, controls the quality of the HOP metric's routes. For example, originally LQSR sent Link Info messages when using HOP. When we changed that, to make LQSR with HOP behave more like DSR, we saw a significant improvement in median TCP throughput. This is because the background load of Link Info messages was repopulating the link caches too quickly, reducing the effectiveness of the Route Error messages.

Our study has several limitations that we would like to correct in future work. First, our data traffic is entirely artificial. We would prefer to measure the performance of real network traffic generated by real users. Second, we do not investigate packet loss and jitter with constant-bit-rate datagram traffic. This would be relevant to the performance of multimedia traffic. We would also like to investigate performance of other wireless link quality metrics such as signal strength.

6. A MOBILE SCENARIO

In the traffic scenarios that we have considered so far, all the nodes have been stationary. In community networks like [6, 27, 26] most nodes are indeed likely to be stationary. However, in most other ad hoc wireless networks, at least some of the nodes are mobile. Here, we consider a scenario that involves a single mobile node, and compare the performance of ETX and HOP metrics.

The layout of our testbed is shown in Figure 2. We set up a TCP receiver on node 100. We slowly and steadily walked

around the periphery of the network with a Dell Latitude Laptop, equipped with a NetGear card. A process running on this laptop repeatedly established a TCP connection to the receiver running on node 100, and transferred as much data as it could in 1 minute. We did 15 such transfers in each walk-about. We did 3 such walk-about each for ETX and HOP. Thus, for each metric we did 45 TCP transfers.

The median throughput of these 45 transfers, along with SIQR is shown in Figure 21. We choose median over mean since the distribution of throughputs is highly skewed. The median throughput under HOP metric is 36% higher than the median throughput under the ETX metric. Note also that the SIQR for ETX is 173, which is comparable to the SIQR of 188 for HOP. Since the median throughput under ETX is lower, the higher SIQR indicates greater variability in throughput under ETX.

As the sender moves around the network, the ETX metric does not react sufficiently quickly to track the changes in link quality. As a result, the node tries to route its packets using stale, and sometimes incorrect information. The salvaging mechanisms built into LQSR do help to some extent, but not well enough to overcome the problem completely. Our results with PktPair (not included here) indicate that that this problem is not limited to just the ETX metric. Any approach that tries to measure link quality will need some time to come up with a stable measure of link quality. If during this time the mobile user moves sufficiently, the link quality measurements would not be correct. Note that we do have penalty mechanisms built into our link quality measurements. If a data packet is dropped on a link, we penalize the metric as described in Section 2. We are investigating the possibility that by penalizing the metric more aggressively on data packet drops we can improve the performance of ETX.

The HOP metric, on the other hand, faces no such problems. It uses new links as quickly as the node discovers them. The efficacy of various DSR mechanisms to improve performance in a mobile environment has been well documented [15]. The metric also removes from link cache any link on which a node suffers even one packet loss. This mechanism, which hurts the performance of HOP metric under heavy load, benefits it in the mobile scenario.

We stress that this experiment is not an attempt to draw general conclusions about the suitability of any metric for routing in mobile wireless networks. Indeed, such conclusions can not be drawn from results of a single experiment. This experiment only serves to underscore the fact that static and mobile wireless networks can present two very different sets of challenges, and solutions that work well in one setting are not guaranteed to work just as well in another.

7. RELATED WORK

There is a large body literature comparing the performance of various ad hoc routing protocols. Most of this work is simulation-based and the ad hoc routing protocols studied all minimize hop-count. Furthermore, many of these studies focus on scenarios that involve significant node mobility. For example, Broch et al. [7] compared the performance of DSDV [23], TORA [22], DSR [15], and AODV [24] via simulations.

The problem of devising a link-quality metric for static 80.211 ad hoc networks has been studied previously. Most

notably, De Couto et al. [9] propose ETX and compare its performance to HOP using DSDV and DSR with a small-datagram workload. Their study differs from ours in many aspects. They conclude that ETX outperforms HOP with DSDV, but find little benefit with DSR. They only study the throughput of single, short (30 second) data transfers using small datagrams. Their experiments include no mobility. In contrast, we study TCP transfers. We examine the impact of multiple simultaneous data transfers. We study variable-length data transfers and in particular, look at web-like workloads where latency is more important than throughput. Finally, our work includes a scenario with some mobility. Their implementation of DSR differs from ours in several ways, which may partly explain our different results. They do not have Metric Maintenance mechanisms. In their testbed (as in ours), the availability of multiple paths means after the initial route discovery the nodes rarely send Route Requests. Hence during their experiments, the sender effectively routes using a snapshot of the ETX metrics discovered at the start of the experiment. Their implementation takes advantage of 802.11 link-layer acknowledgments for failure detection. This means their link-failure detection is not vulnerable to loss, or perceived loss due to delay. Their implementation does not support salvaging. They mitigate this in their experiments by sending five “priming” packets before starting each data transfer. Their implementation uses a “blacklist” mechanism to cope with asymmetric links. Finally, their implementation has no security support and does not use cryptography so it has much less CPU overhead.

Woo et al. [28] examines the interaction of link quality and ad hoc routing for sensor networks. Their scheme is based on passive observation of packet reception probability. Using this probability, they compare several routing protocols including shortest-path routing with thresholding to eliminate links with poor quality and ETX-based distance-vector routing. Their study uses rate-limited datagram traffic. They conclude that ETX-based routing is more robust.

Signal-to-noise ratio (SNR), has been used as a link quality metric in several routing schemes for mobile ad hoc networks. For example, in [12] the authors use an SNR threshold value to filter links discovered by DSR Route Discovery. The main problem with these schemes is that they may end up excluding links that are necessary to maintain connectivity. Another approach is used in [11], where links are still classified as good and bad based on a threshold value, but a path is permitted to use poor-quality links to maintain connectivity. Punnoose et. al. [25] also use signal strength as a link quality metric. They convert the predicted signal strength into a link quality factor, which is used as assign weights to the links. Zhao and Govindan [29]. studied packet delivery performance in sensor networks, and discovered that high signal strength implies low packet loss, but low signal strength does not imply high packet loss. We plan to study the SNR metric in our testbed as part of our future work. Our existing hardware and software setup does not provide adequate support to study this metric.

Awerbuch et. al. [4] study impact of automatic rate selection on performance of ad hoc networks. They propose a routing algorithm that selects a path with minimum transmission time. Their metric does not take packet loss into account. It is possible to combine this metric with the ETX metric, and study performance of the combined metric. This is also part of our future work.

An implementation of AODV that uses the link-filtering approach, based on measurement of loss rate of unicast probes, was demonstrated in a recent IETF meeting [3, 19]. We plan to test this implementation in our testbed.

8. CONCLUSIONS

We have examined the performance of three candidate link-quality metrics for ad hoc routing and compared them to minimum hop-count routing. Our results are based on several months of experiments using a 23-node static ad hoc network in an office environment. The results show that with stationary nodes the ETX metric significantly outperforms hop-count. The RTT and PktPair metrics perform poorly because they are load-sensitive and hence suffer from self-interference. However, in a mobile scenario hop-count performs better because it reacts more quickly to fast topology change.

Acknowledgments

Yih-Chun Hu implemented DSR within the MCL framework as part of his internship project. This was our starting point for developing LQSR. We would like to thank Atul Adya, Victor Bahl and Alec Wolman for several helpful discussions and suggestions. We would also like to thank the anonymous reviewers for their feedback. Finally, we would like to thank the support staff at Microsoft Research for their help with various system administration issues.

9. REFERENCES

- [1] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for IEEE 802.11 wireless networks. In *BroadNets*, 2004.
- [2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In *SOSP*, 2001.
- [3] AODV@IETF. <http://moment.cs.ucsb.edu/aodv-ietf/>.
- [4] B. Awerbuch, D. Holmer, and H. Rubens. High throughput route selection in multi-rate ad hoc wireless networks. Technical report, Johns Hopkins CS Dept, March 2003. v 2.
- [5] P. Bardford and M. Crovella. Generating representative web workloads for network and server performance evaluation. In *SIGMETRICS*, Nov. 1998.
- [6] Bay area wireless users group. <http://www.bawug.org/>.
- [7] J. Broch, D. Maltz, D. Johnson, Y.-C. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *MOBICOM*, Oct. 1998.
- [8] D. De Couto. Personal communication, Nov. 2003.
- [9] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. High-throughput path metric for multi-hop wireless routing. In *MOBICOM*, Sep 2003.
- [10] R. Draves, J. Padhye, and B. Zill. The architecture of the Link Quality Source Routing Protocol. Technical Report MSR-TR-2004-57, Microsoft Research, 2004.
- [11] T. Goff, N. Abu-Ahazaleh, D. Phatak, and R. Kahvecioglu. Preemptive routing in ad hoc networks. In *MOBICOM*, 2001.
- [12] Y.-C. Hu and D. B. Johnson. Design and demonstration of live audio and video over multi-hop wireless networks. In *MILCOM*, 2002.
- [13] P. Huang and J. Heidemann. Capturing tcp burstiness for lightweight simulation. In *SCS Multiconference on Distributed Simulation*, Jan. 2001.
- [14] R. Jain. *The Art of Computer Systems Performance Analysis*. John Wiley and Sons, Inc., 1991.
- [15] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad-hoc wireless networks. In T. Imielinski and H. Korth, editors, *Mobile Computing*. Kluwer Academic Publishers, 1996.
- [16] R. Karrer, A. Sabharwal, and E. Knightly. Enabling Large-scale Wireless Broadband: The Case for TAPs. In *HotNets*, Nov 2003.
- [17] S. Keshav. A Control-theoretic approach to flow control. In *SIGCOMM*, Sep 1991.
- [18] A. Khanna and J. Zinky. The Revised ARPANET Routing Metric. In *SIGCOMM*, 1989.
- [19] L. Krishnamurthy. Personal communication, Dec. 2003.
- [20] J. Moy. OSPF Version 2. RFC2328, April 1998.
- [21] K. Park, G. Kim, and M. Crovella. On the relationship between file sizes, transport protocols and self-similar network traffic. In *ICNP*, 1996.
- [22] V. D. Park and M. S. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *INFOCOM*, Apr 1997.
- [23] C. E. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance vector routing (dssdv) for mobile computers. In *SIGCOMM*, Sep. 1994.
- [24] C. E. Perkins and E. M. Royer. Ad-hoc on-demand distance vector routing. In *WMCSA*, Feb 1999.
- [25] R. Punnoose, P. Nitkin, J. Borch, and D. Stancil. Optimizing wireless network protocols using real time predictive propagation modeling. In *RAWCON*, Aug 1999.
- [26] MIT roofnet. <http://www.pdos.lcs.mit.edu/roofnet/>.
- [27] Seattle wireless. <http://www.seattlewireless.net/>.
- [28] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *SenSys*, Nov 2003.
- [29] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *SenSys*, Nov. 2003.